

Easy Monotonic Policy Iteration

Joshua Achiam
UC Berkeley

JACHIAM@BERKELEY.EDU

Abstract

A key problem in reinforcement learning for control with general function approximators (such as deep neural networks and other nonlinear functions) is that, for many algorithms employed in practice, updates to the policy or Q -function may fail to improve performance—or worse, actually cause the policy performance to degrade. Prior work has addressed this for policy iteration by deriving tight policy improvement bounds; by optimizing the lower bound on policy improvement, a better policy is guaranteed. However, existing approaches suffer from bounds that are hard to optimize in practice because they include sup norm terms which cannot be efficiently estimated or differentiated. In this work, we derive a better policy improvement bound where the sup norm of the policy divergence has been replaced with an average divergence; this leads to an algorithm, Easy Monotonic Policy Iteration, that generates sequences of policies with guaranteed non-decreasing returns and is easy to implement in a sample-based framework.

1. Introduction

Following the success of the Deep Q-Network (DQN) approach (Mnih et al., 2013), there has been a surge of interest in using reinforcement learning for control with nonlinear function approximators, particularly with deep neural networks; in these methods, the policy or the Q -function is represented by the approximator. Examples include variants on DQN, such as Double-DQN (van Hasselt et al., 2015) and Deep Recurrent Q-Learning (Hausknecht & Stone, 2015), as well as methods that mix neural network policies and neural network value functions, such as the asynchronous advantage actor-critic algorithm (Mnih et al., 2016). However, despite empirical successes, there are few algorithms that come with theoretical guarantees for continued policy improvement when training policies represented by arbitrary nonlinear function approximators.

One approach, which serves as the inspiration for this work, seeks to maximize a lower bound on policy performance to guarantee an improvement. This method has its roots in conservative policy iteration (CPI) (Kakade & Langford, 2002), and was extended separately by Pirotta et al. (2013) and Schulman et al. (2015a). Both Pirotta et al. and Schulman et al. derived similar policy improvement bounds consisting of two parts: an expected advantage of the new policy with respect to the old policy, and a penalty on a divergence between the new policy and the old policy. The divergence penalty, in both cases, is quite steep: it involves the maximum policy divergence over all states. This makes it particularly difficult to apply the bounds in the usual situations where function approximation is desirable: domains where the model—and hence the total set of states—is unknown (in which case it is not possible to evaluate the bound) and/or where the state space is large (in which case the bound may be unnecessarily conservative). Pirotta et al. developed algorithms that primarily apply to the case where the model is known or where approximation is present only in the advantage estimation, and so did not address this issue. Schulman et al. addressed the issue by proposing to solve an approximate form of the problem, where the maximum divergence penalty is replaced by a trust-region constraint on the average divergence; this algorithm is called trust region policy optimization (TRPO). They found TRPO worked quite well in a number of domains, successfully optimizing neural network policies to play Atari from raw pixels and to control a simulated robot in locomotion tasks.

In this work, we derive a new policy improvement bound where the penalty on policy divergence goes as an average, rather than the maximum, divergence. This allows us to propose Easy Monotonic Policy Iteration (EMPI), an algorithm that exploits the bound to generate sequences of policies with guaranteed non-decreasing returns, which is easy to implement in a sample-based framework. It also enables us to give a new theoretical justification for TRPO: we are able to show that each iteration of TRPO has a worst-case degradation of policy performance which depends on a hyperparameter of the algorithm. Our contributions at present are entirely theoretical, but empirical results from testing EMPI will appear in a future version of this work.

2. Preliminaries

A Markov decision process is a tuple, (S, A, R, P, μ) , where S is the set of states, A is the set of actions, $R : S \times A \times S \rightarrow \mathbb{R}$ is the reward function, $P : S \times A \times S \rightarrow [0, 1]$ is the transition probability function (where $P(s'|s, a)$ is the probability of transitioning to state s' given that the previous state was s and the agent took action a in s), and $\mu : S \rightarrow [0, 1]$ is the starting state distribution. A policy $\pi : S \times A \rightarrow [0, 1]$ is a distribution over actions per state, with $\pi(a|s)$ the probability of selecting a in state s . We consider the problem of picking a policy π that maximizes the expected infinite-horizon discounted total reward,

$$J(\pi) \doteq \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) \right],$$

where $\gamma \in [0, 1]$ is the discount factor, τ denotes a trajectory ($\tau = (s_0, a_0, s_1, \dots)$), and $\tau \sim \pi$ is shorthand for indicating that trajectories are drawn from distributions induced by π : $s_0 \sim \mu$, $a_t \sim \pi(\cdot|s_t)$, $s_{t+1} \sim P(\cdot|s_t, a_t)$.

We define the on-policy value function, $V^\pi : S \rightarrow \mathbb{R}$, action-value function $Q^\pi : S \times A \rightarrow \mathbb{R}$, and advantage function $A^\pi : S \times A \rightarrow \mathbb{R}$ in the usual way:

$$\begin{aligned} V^\pi(s) &\doteq \mathbb{E}_{a_0, s_1, \dots} \left[\sum_{t=0}^{\infty} \gamma^t R_t \middle| s_0 = s \right], \\ Q^\pi(s, a) &\doteq \mathbb{E}_{s_1, a_1, \dots} \left[\sum_{t=0}^{\infty} \gamma^t R_t \middle| s_0 = s, a_0 = a \right], \\ A^\pi(s, a) &\doteq Q^\pi(s, a) - V^\pi(s), \end{aligned}$$

where $R_t = R(s_t, a_t, s_{t+1})$. The Q and V functions are connected by

$$Q^\pi(s, a) = \mathbb{E}_{s' \sim P(\cdot|s, a)} [R(s, a, s') + \gamma V^\pi(s') | s, a].$$

Our analysis will make extensive use of the discounted future state distribution, d^π , which is defined as

$$d^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi).$$

It allows us to express the expected discounted total reward compactly as

$$J(\pi) = \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [R(s, a, s')], \quad (1)$$

where by $a \sim \pi$, we mean $a \sim \pi(\cdot|s)$, and by $s' \sim P$, we mean $s' \sim P(\cdot|s, a)$. We drop the explicit notation for the sake of reducing clutter, but it should be clear from context that a and s' depend on s .

Next, we will examine some useful properties of d^π that become apparent in vector form for finite state spaces. Let $p_\pi^t \in \mathbb{R}^{|S|}$ denote the vector with components $p_\pi^t(s) = P(s_t = s | \pi)$, and let $P_\pi \in \mathbb{R}^{|S| \times |S|}$ denote the transition matrix with components $P_\pi(s'|s) = \int da P(s'|s, a) \pi(a|s)$; then $p_\pi^t = P_\pi p_\pi^{t-1} = P_\pi^t \mu$ and

$$\begin{aligned} d^\pi &= (1 - \gamma) \sum_{t=0}^{\infty} (\gamma P_\pi)^t \mu \\ &= (1 - \gamma)(I - \gamma P_\pi)^{-1} \mu. \end{aligned} \quad (2)$$

This formulation helps us easily obtain the following lemma.

Lemma 1. For any function $f : S \rightarrow \mathbb{R}$ and any policy π ,

$$(1 - \gamma) \mathbb{E}_{s \sim \mu} [f(s)] + \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [\gamma f(s')] - \mathbb{E}_{s \sim d^\pi} [f(s)] = 0. \quad (3)$$

Proof. Multiply both sides of (2) by $(I - \gamma P_\pi)$ and take the inner product with the vector $f \in \mathbb{R}^{|S|}$. \square

Combining this with (1), we obtain the following, for any function f and any policy π :

$$\begin{aligned} J(\pi) &= \mathbb{E}_{s \sim \mu} [f(s)] \\ &\quad + \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [R(s, a, s') + \gamma f(s') - f(s)]. \end{aligned} \quad (4)$$

This identity is nice for two reasons. First: if we pick f to be an approximator of the value function V^π , then (4) relates the true discounted return of the policy ($J(\pi)$) to the estimate of the policy return ($\mathbb{E}_{s \sim \mu} [f(s)]$) and to the on-policy average TD-error of the approximator; this is aesthetically satisfying. Second: it shows that reward-shaping by $\gamma f(s') - f(s)$ has the effect of translating the total discounted return by $\mathbb{E}_{s \sim \mu} [f(s)]$, a fixed constant independent of policy; this illustrates the finding of Ng. et al. (1999) that reward shaping by $\gamma f(s') + f(s)$ does not change the optimal policy.

It is also helpful to introduce an identity for the vector difference of the discounted future state visitation distributions on two different policies, π' and π . Define the matrices $G \doteq (I - \gamma P_\pi)^{-1}$, $\bar{G} \doteq (I - \gamma P_{\pi'})^{-1}$, and $\Delta = P_{\pi'} - P_\pi$. Then:

$$\begin{aligned} G^{-1} - \bar{G}^{-1} &= (I - \gamma P_\pi) - (I - \gamma P_{\pi'}) \\ &= \gamma \Delta; \end{aligned}$$

left-multiplying by G and right-multiplying by \bar{G} , we obtain

$$\bar{G} - G = \gamma \bar{G} \Delta G.$$

Thus

$$\begin{aligned} d^{\pi'} - d^{\pi} &= (1 - \gamma) (\bar{G} - G) \mu \\ &= \gamma(1 - \gamma) \bar{G} \Delta G \mu \\ &= \gamma \bar{G} \Delta d^{\pi}. \end{aligned} \quad (5)$$

For simplicity in what follows, we will only consider MDPs with finite state and action spaces, although our attention is on MDPs that are too large for tabular methods.

3. Main Results

In this section, we will derive and present the new policy improvement bound. We will begin with a lemma:

Lemma 2. *For any function $f : S \rightarrow \mathbb{R}$ and any policies π' and π , define*

$$\begin{aligned} L_{\pi, f}(\pi') &\doteq \\ &\mathbb{E}_{\substack{s \sim d^{\pi} \\ a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) (R(s, a, s') + \gamma f(s') - f(s)) \right], \end{aligned} \quad (6)$$

and $\epsilon_f^{\pi'} \doteq \max_s |\mathbb{E}_{a \sim \pi', s' \sim P} [R(s, a, s') + \gamma f(s') - f(s)]|$. Then the following bound holds:

$$J(\pi') - J(\pi) \geq \frac{1}{1 - \gamma} \left(L_{\pi, f}(\pi') - 2\epsilon_f^{\pi'} D_{TV}(d^{\pi'} \| d^{\pi}) \right), \quad (7)$$

where D_{TV} is the total variational divergence. Furthermore, the bound is tight (when $\pi' = \pi$, the LHS and RHS are identically zero).

Proof. First, for notational convenience, let $\delta_f(s, a, s') \doteq R(s, a, s') + \gamma f(s') - f(s)$. (The choice of δ to denote this quantity is intentionally suggestive—this bears a strong resemblance to a TD-error.) By (4), we obtain the identity

$$\begin{aligned} J(\pi') - J(\pi) &= \frac{1}{1 - \gamma} \left(\mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] - \mathbb{E}_{\substack{s \sim d^{\pi} \\ a \sim \pi \\ s' \sim P}} [\delta_f(s, a, s')] \right) \end{aligned}$$

Now, we restrict our attention to the first term in this equation. Let $\bar{\delta}_f^{\pi'} \in \mathbb{R}^{|S|}$ denote the vector of components $\bar{\delta}_f^{\pi'}(s) = \mathbb{E}_{a \sim \pi', s' \sim P} [\delta_f(s, a, s') | s]$. Observe that

$$\begin{aligned} \mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] &= \langle d^{\pi'}, \bar{\delta}_f^{\pi'} \rangle \\ &= \langle d^{\pi}, \bar{\delta}_f^{\pi'} \rangle + \langle d^{\pi'} - d^{\pi}, \bar{\delta}_f^{\pi'} \rangle \end{aligned}$$

This term is then straightforwardly bounded by applying Hölder's inequality; for any $p, q \in [1, \infty]$ such that $1/p + 1/q = 1$, we have

$$\mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] \geq \langle d^{\pi}, \bar{\delta}_f^{\pi'} \rangle - \|d^{\pi'} - d^{\pi}\|_p \|\bar{\delta}_f^{\pi'}\|_q.$$

Particularly, we choose $p = 1$ and $q = \infty$; however, we believe that this step is very interesting, and different choices for dealing with the inner product $\langle d^{\pi'} - d^{\pi}, \bar{\delta}_f^{\pi'} \rangle$ may lead to novel and useful bounds.

With $\|d^{\pi'} - d^{\pi}\|_1 = 2D_{TV}(d^{\pi'} \| d^{\pi})$ and $\|\bar{\delta}_f^{\pi'}\|_{\infty} = \epsilon_f^{\pi'}$, the bound is almost obtained. The last step is to observe that, by the importance sampling identity,

$$\begin{aligned} \langle d^{\pi}, \bar{\delta}_f^{\pi'} \rangle &= \mathbb{E}_{\substack{s \sim d^{\pi} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] \\ &= \mathbb{E}_{\substack{s \sim d^{\pi} \\ a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} \right) \delta_f(s, a, s') \right]. \end{aligned}$$

After grouping terms, the bound is obtained. \square

This lemma makes use of many ideas that have been explored before; for the special case of $f = V^{\pi}$, this strategy (after bounding $D_{TV}(d^{\pi'} \| d^{\pi})$) leads directly to some of the policy improvement bounds previously obtained by Pirota et al. and Schulman et al. The form given here is more general, however, because it allows for freedom in choosing f . (Although we do not report the results here, we note that this freedom allows for the derivation of analogous bounds involving Bellman errors of Q-function approximators, which is interesting and suggestive.)

Remark. It is reasonable to ask if there is a choice of f which maximizes the lower bound here. This turns out to trivially be $f = V^{\pi'}$. Observe that $\mathbb{E}_{s' \sim P} [\delta_{V^{\pi'}}(s, a, s') | s, a] = A^{\pi'}(s, a)$. For all states, $\mathbb{E}_{a \sim \pi'} [A^{\pi'}(s, a)] = 0$ (by the definition of $A^{\pi'}$), thus $\bar{\delta}_{V^{\pi'}}^{\pi'} = 0$ and $\epsilon_{V^{\pi'}}^{\pi'} = 0$. Also, $L_{\pi, V^{\pi'}}(\pi') = -\mathbb{E}_{s \sim d^{\pi}, a \sim \pi} [A^{\pi'}(s, a)]$; from (4) with $f = V^{\pi'}$, we can see that this exactly equals $J(\pi') - J(\pi)$. Thus, for $f = V^{\pi'}$, we recover an exact equality. While this is not practically useful to us (because, when we want to optimize a lower bound with respect to π' , it is too expensive to evaluate $V^{\pi'}$ for each candidate to be practical), it provides insight: the penalty coefficient on the divergence captures information about the mismatch between f and $V^{\pi'}$.

Next, we are interested in bounding the divergence term, $\|d^{\pi'} - d^{\pi}\|_1$. We give the following lemma; to the best of our knowledge, this is a new result.

Lemma 3. *The divergence between discounted future state visitation distributions, $\|d^{\pi'} - d^\pi\|_1$, is bounded by an average divergence of the policies π' and π :*

$$\|d^{\pi'} - d^\pi\|_1 \leq \frac{2\gamma}{1 - \gamma} \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]], \quad (8)$$

where $D_{TV}(\pi' || \pi)[s] = (1/2) \sum_a |\pi'(a|s) - \pi(a|s)|$.

Proof. First, using (5), we obtain

$$\begin{aligned} \|d^{\pi'} - d^\pi\|_1 &= \gamma \|\bar{G} \Delta d^\pi\|_1 \\ &\leq \gamma \|\bar{G}\|_1 \|\Delta d^\pi\|_1. \end{aligned}$$

$\|\bar{G}\|_1$ is bounded by:

$$\|\bar{G}\|_1 = \|(I - \gamma P_{\pi'})^{-1}\|_1 \leq \sum_{t=0}^{\infty} \gamma^t \|P_{\pi'}\|_1^t = (1 - \gamma)^{-1}$$

To conclude the lemma, we bound $\|\Delta d^\pi\|_1$.

$$\begin{aligned} \|\Delta d^\pi\|_1 &= \sum_{s'} \left| \sum_s \Delta(s'|s) d^\pi(s) \right| \\ &\leq \sum_{s, s'} |\Delta(s'|s)| d^\pi(s) \\ &= \sum_{s, s'} \left| \sum_a P(s'|s, a) (\pi'(a|s) - \pi(a|s)) \right| d^\pi(s) \\ &\leq \sum_{s, a, s'} P(s'|s, a) |\pi'(a|s) - \pi(a|s)| d^\pi(s) \\ &= \sum_{s, a} |\pi'(a|s) - \pi(a|s)| d^\pi(s) \\ &= 2 \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]]. \end{aligned}$$

□

The new policy improvement bound follows immediately.

Theorem 1. *For any function $f : S \rightarrow \mathbb{R}$ and any policies π', π , with $L_{\pi, f}(\pi')$ as defined in (6) and $\epsilon_f^{\pi'} \doteq \max_s |\mathbb{E}_{a \sim \pi', s' \sim P}[R(s, a, s') + \gamma f(s') - f(s)]|$, the following bound holds:*

$$\begin{aligned} J(\pi') - J(\pi) &\geq \frac{1}{1 - \gamma} \left(L_{\pi, f}(\pi') - \frac{2\gamma\epsilon_f^{\pi'}}{1 - \gamma} \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]] \right). \end{aligned} \quad (9)$$

Furthermore, the bound is tight (when $\pi' = \pi$, the LHS and RHS are identically zero).

Proof. Begin with the bound of lemma 2 and bound the divergence $D_{TV}(d^{\pi'} || d^\pi)$ by lemma 3. □

A few quick observations connect this result to prior work. Clearly, we could bound the expectation $\mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]]$ by $\max_s D_{TV}(\pi' || \pi)[s]$. Doing this, picking $f = V^\pi$, and bounding $\epsilon_f^{\pi'}$ to get a second factor of $\max_s D_{TV}(\pi' || \pi)[s]$, we recover (up to assumption-dependent factors) the policy improvement bounds given by Pirootta et al. (2013) as Corollary 3.6, and by Schulman et al. (2015a) as Theorem 1a.

Because the choice of $f = V^\pi$ does allow for a nice simplification, we will give the bound with this choice as a corollary.

Corollary 1. *For any policies π', π , with $\epsilon^{\pi'} \doteq \max_s |\mathbb{E}_{a \sim \pi'} [A^\pi(s, a)]|$, the following bound holds:*

$$\begin{aligned} J(\pi') - J(\pi) &\geq \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} \left[A^\pi(s, a) - \frac{2\gamma\epsilon^{\pi'}}{1 - \gamma} D_{TV}(\pi' || \pi)[s] \right]. \end{aligned} \quad (10)$$

4. Easy Monotonic Policy Iteration

As with the weaker versions of the bound, this bound allows us to generate a sequence of policies with non-degrading performance in any restricted class of policies, i.e. policies that are represented by neural networks or other function approximators. We'll use Π_θ to denote an arbitrary restricted class of policies, and understand that this usually means policies smoothly parametrized by some set of parameters θ . Algorithm 1 gives the general template for Easy Monotonic Policy Iteration (EMPI), which obtains monotonic improvements using (9) (and one additional small step of bounding to remove π' from the penalty coefficient). To see that EMPI is indeed monotonic, observe that π_i is a feasible point of the optimization defining π_{i+1} , with objective value 0; this is a certificate that the objective at optimum is ≥ 0 .

Although we do not specify how to solve the optimization problem, we note that the objective is differentiable with respect to the parameters of a candidate policy π' : as a result, a gradient-based method can be used. Furthermore, when we use neural network policies where the vector of n parameters may take on any value in \mathbb{R}^n , the optimization is unconstrained. This is one sense in which we consider this algorithm 'easy' to implement.

Usually we would be interested in applying this method to problems with large state or unknown state spaces, where exact calculation of the objective function is not feasible. Because the objective is defined almost entirely in terms of expectations on policy π_i , however, we can use a sample-based approximation of it; this is the other sense in which this algorithm is 'easy.' The only challenge is estimat-

Algorithm 1 Easy Monotonic Policy Iteration: monotonic policy improvements in arbitrary policy classes

Input: Initial policy $\pi_0 \in \Pi_\theta$, max number of iterations N , stopping tolerance α
for $i = 1, 2, \dots, N$ **or until** $J(\pi_i) - J(\pi_{i-1}) \leq \alpha$ **do**
 Choose function $f_i : S \rightarrow \mathbb{R}$
 $\pi_{i+1} \leftarrow \arg \max_{\pi' \in \Pi_\theta} L_{\pi_i, f_i}(\pi') - C \mathbb{E}_{s \sim d^{\pi_i}} [D_{TV}(\pi' || \pi_i)]$,
 where $C = \frac{2\gamma\epsilon}{1-\gamma}$,
 and $\epsilon = \max_{s,a} |\mathbb{E}_{s' \sim P}[R(s, a, s') + \gamma f_i(s') - f_i(s)]|$.
end for

ing ϵ , where we might apply a worst-case bound (potentially making the policy step too conservative) or a reasonable heuristic bound (potentially permitting degraded policy performance).

4.1. Implementing EMPI

A practical form of the general algorithm is obtained by choosing a particular form for the reward-shaping functions f_i , possibly approximating terms in the objective, and replacing the expectations in the objective by sample estimates. We'll choose $f_i = V^{\pi_i}$, so that the base optimization problem becomes

$$\max_{\pi' \in \Pi_\theta} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} \left[A^\pi(s, a) - \frac{2\gamma\epsilon}{1-\gamma} D_{TV}(\pi' || \pi)[s] \right],$$

where $\epsilon = \max_{s,a} |A^\pi(s, a)|$.

Suppose that we use an estimator of the advantage, $\hat{A}^\pi(s, a)$, instead of the true advantage $A^\pi(s, a)$. For example, if we have learned a neural network value function approximator \hat{V}^π , we may use $\hat{A}^\pi(s, a) = R(s, a, s') + \gamma \hat{V}^\pi(s') - \hat{V}^\pi(s)$; or perhaps we might use the generalized advantage estimator (Schulman et al., 2015b). Observe that, because $\mathbb{E}_{a \sim \pi}[A^\pi(s, a)] = 0$, for every state s we have

$$\begin{aligned} \mathbb{E}_{a \sim \pi'} [A^\pi(s, a)] &= \\ & \mathbb{E}_{a \sim \pi'} [\hat{A}^\pi(s, a)] - \mathbb{E}_{a \sim \pi} [\hat{A}^\pi(s, a)] \\ & + \sum_a (\pi'(a|s) - \pi(a|s)) (A^\pi(s, a) - \hat{A}^\pi(s, a)). \end{aligned}$$

From this, we derive a bound:

$$\begin{aligned} \mathbb{E}_{a \sim \pi'} [A^\pi(s, a)] &\geq \\ & \mathbb{E}_{a \sim \pi'} [\hat{A}^\pi(s, a)] - \mathbb{E}_{a \sim \pi} [\hat{A}^\pi(s, a)] \\ & - 2 \max_a |A^\pi(s, a) - \hat{A}^\pi(s, a)| D_{TV}(\pi' || \pi)[s], \end{aligned}$$

which gives us the following corollary to Theorem 1.

Corollary 2 (Policy Improvement Bound with Arbitrary Advantage Estimators). *For any policies π', π , and any advantage estimator $\hat{A}^\pi : S \times A \rightarrow \mathbb{R}$, with $c(s) \doteq \max_a |A^\pi(s, a) - \hat{A}^\pi(s, a)|$ and $\epsilon^{\pi'} \doteq \max_s |\mathbb{E}_{a \sim \pi'} [A^\pi(s, a)]|$, the following bound holds:*

$$\begin{aligned} J(\pi') - J(\pi) &\geq \frac{1}{1-\gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} \left[\hat{A}^\pi(s, a) - \mathbb{E}_{\bar{a} \sim \pi} [\hat{A}^\pi(s, \bar{a})] \right. \\ &\quad \left. - 2 \left(c(s) + \frac{\gamma\epsilon^{\pi'}}{1-\gamma} \right) D_{TV}(\pi' || \pi)[s] \right]. \end{aligned} \quad (11)$$

Furthermore, the bound is tight (when $\pi' = \pi$, the LHS and RHS are identically zero).

Corollary 2 tells us that we are theoretically justified in using any advantage estimators as long as we increase the penalty on the policy divergence appropriately. Also, we can see that if the estimator is high-quality ($|A^\pi(s, a) - \hat{A}^\pi(s, a)|$ small) we can take larger policy improvement steps, which makes sense. If the advantage estimator is poor, then we probably will not make much progress; this is reflected in the bound.

The base optimization problem for EMPI with an advantage estimator, which we obtain by using (11), dropping constants, and removing π' from the penalty coefficient by bounding, is

$$\begin{aligned} \max_{\pi' \in \Pi_\theta} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} & \left[\hat{A}^\pi(s, a) - C(s) D_{TV}(\pi' || \pi)[s] \right], \\ \text{where } C(s) &= 2 \left(c(s) + \frac{\gamma\epsilon}{1-\gamma} \right), \\ c(s) &= \max_a |A^\pi(s, a) - \hat{A}^\pi(s, a)|, \\ \text{and } \epsilon &= \max_{s,a} |A^\pi(s, a)|. \end{aligned}$$

Now, we will put the objective in a form which can be sampled on policy π . First, we use the importance sampling identity so that the objective becomes

$$\mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi}} \left[\frac{\pi'(a|s)}{\pi(a|s)} \hat{A}^\pi(s, a) - C(s) D_{TV}(\pi' || \pi)[s] \right].$$

Next we rewrite the expectation in terms of trajectories:

$$\mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t \left(\frac{\pi'(a_t|s_t)}{\pi(a_t|s_t)} \hat{A}^\pi(s_t, a_t) - C(s_t) D_{TV}(\pi' || \pi)[s_t] \right) \right]. \quad (12)$$

After running an agent on policy π to generate a set of sample trajectories, we can estimate (12) by averaging over

the sample trajectories. The sample estimate of (12) then serves as the objective for the optimization step in EMPI.

We do not give experimental results here, but plan to include them in a future version of this work.

5. Implications for Trust Region Policy Optimization

Schulman et al. proposed Trust Region Policy Optimization, where at every iteration the policy is updated from π to π' by solving the following optimization problem:

$$\begin{aligned} \pi' = \arg \max_{\pi' \in \Pi_\theta} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} [A^\pi(s, a)] \\ \text{subject to } \mathbb{E}_{s \sim d^\pi} [D_{KL}(\pi || \pi')[s]] \leq \delta, \end{aligned} \quad (13)$$

where $D_{KL}(\pi || \pi')[s] = \mathbb{E}_{a \sim \pi} \left[\log \frac{\pi(a|s)}{\pi'(a|s)} \right]$, and the policy divergence limit δ is a hyperparameter of the algorithm.

The KL-divergence and the TV-divergence of arbitrary distributions p, q are related by Pinsker's inequality, $D_{TV}(p || q) \leq \sqrt{D_{KL}(p || q)/2}$; combining this with Jensen's inequality, we obtain the following bound:

$$\begin{aligned} \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]] &\leq \mathbb{E}_{s \sim d^\pi} \left[\sqrt{\frac{1}{2} D_{KL}(\pi || \pi')[s]} \right] \\ &\leq \sqrt{\frac{1}{2} \mathbb{E}_{s \sim d^\pi} [D_{KL}(\pi || \pi')[s]]}. \end{aligned} \quad (14)$$

By this bound and Corollary 1, we have a result for the worst-case performance of TRPO.

Corollary 3 (TRPO worst-case performance). *A lower bound on the policy performance difference between policies π and π' , where π' is given by (13) and $\pi \in \Pi_\theta$, is*

$$J(\pi') - J(\pi) \geq \frac{-\sqrt{2\delta}\gamma\epsilon^{\pi'}}{(1-\gamma)^2}, \quad (15)$$

where $\epsilon^{\pi'} = \max_s |\mathbb{E}_{a \sim \pi'} [A^\pi(s, a)]|$.

Proof. π is a feasible point of the optimization problem with objective value 0, so $\mathbb{E}_{s \sim d^\pi, a \sim \pi'} [A^\pi(s, a)] \geq 0$. The rest follows by Corollary 1 and (14), noting that TRPO bounds the average KL-divergence by δ . \square

By (15), we can see that TRPO is theoretically justified, in the sense that an appropriate selection of the hyperparameter δ can guarantee a worst-case loss.

6. Discussion

In this note, we derived a new policy improvement bound in which an average, rather than a maximum, policy divergence is penalized. We proposed Easy Monotonic Policy

Iteration, an algorithm that exploits the bound to generate sequences of policies with guaranteed non-decreasing returns and which is easy to implement in a sample-based framework. We showed how to implement EMPI and theoretically justified the use of advantage estimators in the optimization in the inner loop. Lastly, we showed that our policy improvement bound gives a new theoretical foundation to Trust Region Policy Optimization, an algorithm for approximate monotonic policy improvements proposed by Schulman et al. (2015a) which was shown to perform well empirically on a wide variety of tasks; here, we were able to bound the worst-case performance at each iteration of TRPO.

In a future version of this work, we will give experimental results from implementing EMPI on a range of reinforcement learning benchmarks, including high-dimensional domains like Atari.

References

Hausknecht, Matthew and Stone, Peter. Deep Recurrent Q-Learning for Partially Observable MDPs. *arXiv preprint arXiv:1507.06527*, 2015.

Kakade, Sham and Langford, John. Approximately Optimal Approximate Reinforcement Learning. *Proceedings of the 19th International Conference on Machine Learning*, pp. 267–274, 2002. URL <http://www.cs.cmu.edu/afs/cs/Web/People/jcl/papers/>

Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Graves, Alex, Antonoglou, Ioannis, Wierstra, Daan, and Riedmiller, Martin. Playing Atari with Deep Reinforcement Learning. *arXiv preprint arXiv:...*, pp. 1–9, 2013. URL <http://arxiv.org/abs/1312.5602>.

Mnih, Volodymyr, Badia, Adrià Puigdomènech, Mirza, Mehdi, Graves, Alex, Lillicrap, Timothy P., Harley, Tim, Silver, David, and Kavukcuoglu, Koray. Asynchronous Methods for Deep Reinforcement Learning. pp. 1–28, 2016. URL <http://arxiv.org/abs/1602.01783>.

Ng, Andrew Y., Harada, Daishi, and Russell, Stuart. Policy invariance under reward transformations : Theory and application to reward shaping. *Sixteenth International Conference on Machine Learning*, 3:278–287, 1999. doi: 10.1.1.48.345.

Pirotta, Matteo, Restelli, Marcello, and Calandriello, Daniele. Safe Policy Iteration. *Proceedings of the 30th International Conference on Machine Learning*, 28, 2013.

Schulman, John, Moritz, Philipp, Jordan, Michael, and

Abbeel, Pieter. Trust Region Policy Optimization. 2015a.

Schulman, John, Moritz, Philipp, Levine, Sergey, Jordan, Michael, and Abbeel, Pieter. High-Dimensional Continuous Control Using Generalized Advantage Estimation. pp. 1–9, 2015b.

van Hasselt, Hado, Guez, Arthur, and Silver, David. Deep Reinforcement Learning with Double Q-learning. *arXiv:1509.06461 [cs]*, 2015. URL <http://arxiv.org/abs/1509.06461.pdf>.